

Express Mailing Label No. EL 813 291 051 US

PATENT APPLICATION
Docket No. 14113.26.1

UNITED STATES PATENT APPLICATION

of

RICHARD S. OHRAN

for

**INCREMENTALLY RESTORING A MASS
STORAGE DEVICE TO A PRIOR STATE**

FOOZZT " GST6200T

WURTMAN, NIDJUGEN & SEELE
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE

BACKGROUND OF THE INVENTION

1. Related Application

This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/257,499, entitled "Methods and Systems for Backing Up and Restoring Computer Data," filed December 21, 2000, which is incorporated herein by reference.

2. The Field of the Invention

The present invention relates to backing up and restoring computer data. More specifically, the present invention relates to systems and methods for minimizing the loss of computer data when restoring computer data that was lost due to the interruption.

3. The Prior State of the Art

With the advent of personal computers, businesses worldwide rely on computer data in performing daily business routines. However, a variety of events such as natural disasters, terrorism, or more mundane events such as computer hardware and/or software failures can occur while a computer is processing. These failures often result in causing the computer data to become corrupt, unreliable, or even lost. The corruption or loss of data, such as customer lists, financial transactions, business documents, business transactions, and so forth, can cause havoc to businesses by resulting in the loss of large investments of time and/or money.

The loss or corruption of computer data is particularly devastating in the world of electronic commerce. The Internet has allowed individuals all over the world to conduct business electronically, thereby resulting in the continual upload of electronic orders. However, all copies of the orders are electronic and thus corruption or loss of the

1 electronic computer data can result in the loss of the business represented by the lost
2 orders.

3 Recognizing the commercial value of reliable computer data, businesses seek
4 ways to protect their data and to reconstruct data that has become corrupt, unreliable, or
5 lost. Traditional approaches of data protection and reconstruction have involved creating
6 a backup copy of the computer data. While it is a simple procedure to preserve a backup
7 copy of an individual file on a floppy disk, the creation of a backup copy becomes more
8 difficult as the amount of data increases.

9 Perhaps one of the simplest approaches to creating a backup copy of a large
10 volume of computer data is to copy the data from a mass storage system to an archival
11 device, such as one or more magnetic tapes. This method stores large amounts of
12 computer data at the expense of immediate access to the data. The magnetic tapes are
13 stored either locally or remotely, and the data is copied from the magnetic tapes to the
14 mass storage system when problems arise with the mass storage system.

15 While the use of an archival device to preserve data loss has the advantage of
16 being relatively simple and inexpensive, it also has severe limitations. One such
17 limitation is the amount of time that prevents user accessibility to the computer data
18 while a backup copy is created and while data is reconstructed. The prevention of user
19 accessibility has traditionally been required to ensure that no data has changed during the
20 process. Because user inaccessibility of data is undesirable, backup copies are created
21 less frequently, thereby causing the computer data in the backup copy to become stale.
22 Similarly, transferring computer data from a magnetic tape to a mass storage system can
23 become very lengthy because the computer data is transferred on a file-by-file basis. The
24 time is further lengthened when the archival mass storage device is remotely located and

FOOTNOTES

1 is not accessible over a network. These long reconstruction periods result in extended
2 computer inaccessibility and cost businesses increased amounts of time and money.

3 Another limitation of the traditional methods is that the backup copy represents
4 data as it existed at a previous instant in time, meaning that the backup copy is not current
5 with subsequent changes made to the original copy. The creation of a backup copy
6 provides the security that a large portion of the computer data can be recovered in the
7 event that the original copy becomes corrupt or lost. This limits the loss to include only
8 the changes made to the original copy since the creation of the last backup copy.
9 However, in some businesses, if the computer data is not current, the data is stale,
10 unreliable, and even useless. This is particularly troubling in the financial world where
11 rates and information change with great frequency. Thus, when the computer data
12 becomes corrupt or lost, the businesses that rely on information that changes frequently
13 are exposed to the risk of losing all of their valuable computer data.



14 It would, therefore, represent an advancement in the art to have a system and
15 method for backing up computer data that could further (minimize the amount of
16 computer data that is lost due to a computer failure or corruption of data) It would also
17 represent an advancement in the art to have a system that allowed data to be backed up
18 without terminating user access to the mass storage system.

SUMMARY OF THE INVENTION

The present invention relates to systems and methods for backing up and restoring computer data. The invention enables computer data to be restored in an incremental manner in the event that data is corrupted or lost. In particular, if data is lost or corrupted, the data can be incrementally advanced through successively older states until a valid set of data is obtained. In this manner, data can be restored to a state that is newer than that associated with a full mirrored or archived copy of the data. (Thus, full mirror or archiving operations on a volume of data can be less frequent without the risk of losing changes to the volume of data that have occurred since the last full mirror or archiving operation.)

According to the invention, a mass storage device stores a plurality of data blocks at time T_0 . At T_0 , a mirrored copy of the data stored of the mass storage device may be made and stored such that the data at T_0 can be conveniently restored, if necessary. Obtaining a mirrored copy of data is often a time or resource consuming process that is preferably conducted relatively infrequently. In order to preserve the data that is changed after the mirrored copy of data is created and before a potential event causing loss of data, data blocks that are to be overwritten after T_0 are stored in a preservation memory. In particular, if, after T_0 , a specified data block is to be overwritten as part of a write

data restored in an incremental manner
— no losing changes of data.



block is stored in a preservation memory prior to the mass storage device. (In addition, the data is time-stamped or otherwise marked so as to or to designate a chronological position of the operations.)

1 As write operations are successively performed after time T_0 , the original data
2 blocks that are to be overwritten are sequentially stored in the preservation memory with
3 an associated time stamp. Thus, the data blocks that are overwritten or otherwise
4 changed after T_0 are preserved in a preservation memory and the time or the order in
5 which the data blocks were overwritten in the mass storage device is specified.

6 In the event that certain data blocks in the mass storage unit device are lost or
7 become corrupted, the data blocks stored in the preservation memory can be used to
8 incrementally restore or reconstruct a valid set of data without reverting completely back
9 to the data as it exists at time T_0 . If, for example, invalid or corrupted data is written to
10 certain data blocks in the mass storage device after time T_0 , the original, valid data blocks
11 are stored in a preservation memory as described above. Using the time stamps
12 specifying the chronological sequence in which the data blocks stored in the preservation
13 memory were overwritten in the mass storage device, the data blocks in the preservation
14 memory are written to the current data stored in the mass storage device.

15 After one or more data blocks from the preservation memory are written back to
16 the mass storage device data in reverse chronological order, a valid set of data is
17 eventually obtained at the mass storage device. Thus, the data blocks stored in the
18 preservation memory are used to reconstruct the data without requiring the data to be
19 reverted completely back to the data as it existed at time T_0 .

20 It is noted that, according to the invention, complete mirror or archive operations
21 performed on the full volume of data stored in the mass storage device can be less
22 frequent than would be otherwise required in the absence of the invention. The data
23 blocks stored on preservation memory are used to restored corrupted data to a state more
24 recent than that associated with the most recent full mirrored copy of the data. It is also

1 noted that storing a sequence of overwritten data blocks in the preservation memory is
2 typically less resource intensive than performing a full mirror operation on the data of the
3 mass storage device. In addition, the data stored in the preservation memory in response
4 to the write request can be a data block having fine granularity as opposed to being an
5 entire file. For instance, a data block stored in a preservation memory response to a write
6 operation in the mass storage device can be data associated with a single sector or another
7 portion or segment of the mass storage device. Storing data in a preservation memory in
8 this manner avoids the need of copying entire files in response to write operations and
9 permits the data to be preserved independent of any file structure associated with the
10 data.

11 Additional features and advantages of the invention will be set forth in the
12 description that follows, and in part will be obvious from the description, or may be
13 learned by the practice of the invention. The features and advantages of the invention
14 may be realized and obtained by means of the instruments and combinations particularly
15 pointed out in the appended claims. These and other features of the present invention
16 will become more fully apparent from the following description and appended claims, or
17 may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the manner in which the above recited and other advantages and features of the invention are obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof that are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a block diagram representing the computer system having an associated mass storage device and preservation memory in which the invention can be implemented.

Figure 2 is a block diagram of the computer system of Figure 1, showing data stored in the mass storage device at time T_0 .

Figure 3 illustrates a sequence of write operations in which data blocks that are to be overwritten are stored in a preservation memory.

Figure 4 illustrates an operation in which data in a mass storage device is incrementally restored in response to a data corruption event.

Figure 5 illustrates the operation of a virtual device that appears as if it contains a copy of data as it existed previously in a mass storage device.

Figure 6 illustrates a computer system representing an example of a suitable operating environment for the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention relates to backing up and restoring computer data. In particular, the invention enables corrupted data to be restored without requiring the data to be reverted completely back to a most recent full-mirrored copy of the data. Data blocks that are to be overwritten after a point in time in which a mirrored copy of the data has been created are stored in a preservation memory. The data blocks are associated with a time stamp other information that designates the time or the chronological order in which the data blocks were overwritten or the original data blocks in the mass storage device were overwritten. In general, the term "time stamp" refers to any such information designating the chronological order or the time of the data blocks stored in the preservation memory. The data blocks in the preservation memory can be used in the event of data corruption to incrementally roll the current, corrupted copy of the data in the mass storage device to a valid set of data.

Figure 1 illustrates a computer system in which the invention can be practiced. The computer system of Figure 1 includes a computer 10, a mass storage device 12, a preservation memory 14, and an I/O device 16. Computer 10 can be any computer or other processing device that manages, generates, stores, or otherwise processes data. For instance, computer 10 can be a conventional personal computer, a server, a special purpose computer, or the like. Figure 6, discussed in greater detail below, sets forth details of one example of a computer in which the invention can be practiced.

Mass storage device 12 is associated with computer 10 and is used to store data obtained from computer 10. In general, mass storage device 12 is a writable, nonvolatile mass storage device. In particular, mass storage device 12 can be the hard drive

1 associated with a conventional personal computer or any other storage volume that is
2 used to store data obtained from computer 10.

3 Preservation memory 14 is a physical or logical device associated with computer
4 10 in which data blocks that are to be overwritten in mass storage device 12 are stored.
5 Preservation memory 14 can be a volatile device, such as a random access memory
6 (RAM) or any other device that can store data blocks that are to be overwritten in mass
7 storage device 12. Although preservation memory 14 is illustrated as being a separate
8 device in Figure 1, the preservation memory can be a partition or another portion of mass
9 storage device 12.

10 As shown in Figure 1, computer system 10 may include an input/output (I/O)
11 device 16 that allows a user to manipulate data processed by computer 10 and stored in
12 mass storage device 12. In general, computer 10 can have I/O device 16 or any other
13 peripheral device that enables users to manipulate data. In addition, computer 10
14 typically includes applications or software having data storage requirements satisfied by
15 the mass storage device 12. In general, however, the invention is broad enough to extend
16 to substantially any computer 10 and associated hardware and software, so long as there
17 is a mass storage device 12 for storing data and preservation memory 14 for storing
18 overwritten data blocks.

19 Figure 2 illustrates the computer system of Figure 1 at a time T_0 , when mass
20 storage device 12 includes a set of data blocks storing given data. In the example of
21 Figure 2, for purposes of illustration, mass storage device 12 is shown as including five
22 data blocks, namely, data blocks A, B, C, D, and E. At T_0 , it is assumed that a mirrored
23 or other backup copy of the full volume of data stored in the mass storage device is
24 created and stored. The mirrored or backup copy of the data obtained at T_0 can be

1 generated in any desired manner, including conventional ways. In one embodiment, the
2 mirrored copy or backup of the data at T_0 can be obtained using the techniques disclosed
3 in U.S. Patent No. 5,649,152, entitled "METHOD AND SYSTEM FOR PROVIDING A
4 STATIC SNAPSHOT OF DATA STORED ON A MASS STORAGE SYSTEM" or U.S.
5 Patent No. 5,835,953, entitled "BACKUP SYSTEM THAT TAKES A SNAPSHOT OF
6 THE LOCATIONS IN A MASS STORAGE DEVICE THAT HAS BEEN IDENTIFIED
7 FOR UPDATING PRIOR TO UPDATING", which are incorporated herein by reference.
8 In general, U.S. Patent No. 5,649,152 discloses a technique for taking a snapshot of data
9 as it exists at a selected moment by preserving data overwritten by write operations. U.
10 S. Patent No. 5,835,953 generally describes techniques for obtaining snapshots or
11 mirrored copies of data by transferring minimal amount of data to a remote location.
12 Alternatively, the mirrored or backup copy of the data of the mass storage device 14 at T_0
13 can be obtained using conventional data transfer processes.

14 Depending on the nature of the data in mass storage device 12, it may be useful to
15 obtain a mirrored or backup copy of the data at T_0 that includes logically consistent data.
16 This is particularly important when the data stored in mass storage device 12 represents
17 transactions, each requiring a sequence of write operations or otherwise requiring a set of
18 I/O operations to be performed in order to have a valid and useful set of data. Data is
19 logically inconsistent when fewer than all of the necessary operations in a transaction or
20 in another required sequence of I/O operations have been performed. Thus, depending on
21 the nature of the data stored in mass storage device 12, the mirrored or backup copy of
22 the data at T_0 may need to be logically consistent data.

23 Obtaining a mirrored or backup copy of the data at T_0 enables the data blocks A,
24 B, C, D, and E stored at mass storage device 12 at T_0 to be reconstructed simply by

1 accessing the mirrored or backup copy of the data. As stated above, however, full mirror
2 or backup operations are often time or bandwidth intensive, such that relatively
3 infrequent complete mirror or backup operations are desirable. As discussed above, and
4 as described in more detail below, the present invention uses preservation memory 14 to
5 store data blocks that enable data of mass storage device 12 to be incrementally restored
6 in reverse chronological order without requiring the data to be reverted completely back
7 to time T_0 . Although periodic mirrored or backup copies of the data of the mass storage
8 device 12 are often useful, the invention can also be practiced solely on the basis of data
9 blocks being stored at preservation memory 14 without periodic mirrored or backup
10 copies.

11 In order to describe the manner in which the data is preserved in the preservation
12 memory after time T_0 , Figure 3 presents a specific example of a set of data blocks in the
13 mass storage device and a sequence of the write operations applied thereto. As shown in
14 Figure 3, the data blocks 20 stored in mass storage device at T_0 are designated as (A, B,
15 C, D, E). At $T_{1.0}$, the computer 10 in Figure 2 issues a write request, whereby data block
16 30 (A_1) is to overwrite existing the data block A, resulting in a set of data blocks 20a (A_1 ,
17 B, C, D, E). In general, in response to a issued write request and prior to the data block
18 in a mass storage device being overwritten with the new data block, the existing data
19 block is stored in the preservation memory, along with a time stamp.

20 In this example, at $T_{1.0}$, data block A is written to the preservation memory, along
21 with a time stamp designating the time $T_{1.0}$ as shown at 40a. Thus, even though data
22 block A has been overwritten in a mass storage device, the data block A is stored in the
23 preservation memory in the event that this data block is needed to reconstruct corrupted
24 data at some future point.

1 At time $T_{1,1}$, data block 30b (D_1) is written to the data stored in the mass storage
2 device, resulting in a set of data blocks 20b (A_1 , B, C, D_1 , E). Prior to the write
3 operation, the data block D, which is to be overwritten, is stored in the preservation
4 memory, along with a time stamp designating the time $T_{1,1}$, as shown at 40b.

5 At time $T_{1,2}$, a data corruption event occurs, resulting in data block 30c (D_x) being
6 written to the data stored in the mass storage device. As used in this example, the
7 subscripts "x" and "y" represent corrupted data. The data corruption can have
8 substantially any cause, such as data entry problems, software problems, hardware
9 problems, or the like. The data corruption event at time $T_{1,2}$ results in the set of data
10 blocks 20c (A_1 , B, C, D_x , E). Prior to the write operation, the data block D_1 , which is to
11 be overwritten, is stored in the preservation memory, along with a time stamp designating
12 the time $T_{1,2}$, as shown at 40c.

13 At time $T_{1,3}$, data block 30d (B_1) is written to the data stored in the mass storage
14 device, resulting in a set of data blocks 20d (A_1 , B_1 , C, D_x , E). Prior to the write
15 operation, the data block B, which is to be overwritten, is stored in the preservation
16 memory, along with a time stamp designating the time $T_{1,3}$, as shown at 40d.

17 Finally, at time $T_{1,4}$, another data corruption event occurs, resulting in data block
18 30e (D_y) being written to the data stored in the mass storage device. The data corruption
19 event at time $T_{1,4}$ results in the set of data blocks 20e (A_1 , B_1 , C, D_y , E). Prior to the write
20 operation, the data block D_x , which is to be overwritten, is stored in the preservation
21 memory, along with a time stamp designating the time $T_{1,4}$, as shown at 40e.

22 Thus, at time $T_{1,4}$, the data blocks 20e existing at the mass storage device are (A_1 ,
23 B_1 , C, D_y , E). As noted above, the data block D_y represents corrupted data, which
24 potentially makes all of the data stored in mass storage device unusable. At $T_{1,4}$,

1 preservation memory also has stored therein data blocks (A, D, D₁, B, D_x) and the
2 corresponding time stamps (T_{1.0}, T_{1.1}, T_{1.2}, T_{1.3}, T_{1.4}).

3 Because the data corruption event has occurred since T₀, a valid set of data can be
4 obtain by reverting completely to the data as it existed in the mass storage device at T₀ by
5 accessing any mirrored or backup copy of the full volume data that was created at T₀.
6 However, reverting completely to the data as it existed at T₀ would result in a loss of all
7 subsequent data written to the mass storage device after T₀. Thus, according to the
8 invention, in response to the data corruption event, the data blocks stored in the
9 preservation memory are used to incrementally restore the data in the mass storage device
10 in reverse chronological order to at a point at which valid, non-corrupted data exists as
11 shown in Figure 4. The preservation memory in Figure 4 at time T₂ includes the set of
12 data blocks and corresponding time stamps that existed in the preservation memory at
13 time T_{1.4}. Likewise, the mass storage device in Figure 4 at time T₂ includes a set of data
14 blocks 20e (A₁, B₁, C, D_y, E) that existed at T_{1.4} in Figure 3.

15 The data restoration operation illustrated in Figure 4 begins with the current set of
16 potentially corrupted data blocks 20e and the data blocks 40e stored in the preservation
17 memory. The time stamps T_{1.0} – T_{1.4} are used to reconstruct the data that previously
18 existed in the mass storage device in reverse chronological order. Thus, the most recently
19 stored data block in preservation memory (D_x) is written to the set of data as it existed at
20 time T₂ to roll back the set of data in the mass storage device to the state in which it
21 existed at T_{1.3}, resulting in a set of data (A₁, B₁, C, D_x, E).

22 It is noted that this operation of rolling back the data of the mass storage device as
23 illustrated in Figure 4 can be performed using the actual data stored in the mass storage
24 device or copy of thereof.

1 Upon rolling back to $T_{1.3}$, it is determined, either by the computer or manually,
2 that the set of data 20d of $T_{1.3}$ still represents corrupted data, in that data block D_x is
3 corrupted. Because the data remains corrupted, the data of the mass storage device is
4 further rolled back to time $T_{1.2}$ using data block B of the preservation memory, which has
5 the time stamp $T_{1.3}$. Thus, the resulting data representing the state of the mass storage
6 device at $T_{1.2}$ is (A_1 , B, C, D_x , E) 20c. Because data blocks 20c also include corrupted
7 data block D_x , it is determined that the set of data blocks 20c represent corrupted data.
8 Accordingly, the data is further rolled back to the state at which it existed at the mass
9 storage device at $T_{1.1}$ by writing to the data at the mass storage device the data block of
10 the preservation memory that is next in reverse chronological order. In particular, data
11 block D_1 having time stamp $T_{1.2}$ is written to the data blocks of mass storage device,
12 resulting in a set of data blocks 20b (A_1 , B, C, D_1 , E), which represents the data as it
13 existed at the mass storage device at $T_{1.1}$.

14 At this point, it is determined that the data 20b (A_1 , B, C, D_1 , E) represents valid,
15 non-corrupted data. Thus, the data blocks of the preservation memory have been used to
16 incrementally restore in reverse chronological order the data blocks of the mass storage
17 device until such time that a valid set of data is obtained. It is also noted that the data
18 blocks 20b (A_1 , B, C, D_1 , E) includes certain data (i.e., A_1 , and D_1) that would have not
19 been included in the restored data had the data been reverted completely back to the
20 mirrored or backup copy data of T_0 . Moreover, this more recent data is restored without
21 requiring a sequence of full mirror or backup operations after T_0 .

22 In view of the foregoing, the operation for restoring the data generally involves
23 applying the data blocks the preservation memory in reverse chronological order to a
24

1 current copy of the data blocks of the mass storage device until such time that the data
2 blocks represent valid non-corrupted data.

3 As can be understood, when the number of write operations is large or frequent,
4 the number of data blocks stored in the preservation memory can increase rapidly. Thus,
5 in practice, there is a trade-off between the frequency of full mirror or backup operations
6 and the volume of data blocks that are stored in the preservation memory. In one
7 embodiment, the frequency at which the full mirror or backup operations are performed is
8 determined by the frequency at which the preservation memory is filled or reaches a
9 certain size. Other words, as the volume of data stored in the preservation memory
10 approaches the capacity of preservation memory, a full mirror or backup operation is
11 performed on the mass storage device. This enables the data in a preservation memory to
12 be discarded, since it is no longer needed in view of the fact that a more recent mirrored
13 or backup copy of the data of the mass storage device has been created.

14 While the invention has been described herein in reference to incrementally
15 restoring corrupted data in reverse chronological order in response to a data corruption
16 event, there are other uses for the basic methods of invention. For instance, the data
17 blocks of the preservation memory can be used to roll the data in the mass storage device
18 back to previous state for other reasons. Indeed, in substantially any situation in which a
19 user wishes to obtain data as it existed previously in a mass storage device.

20 In yet another embodiment, data stored in the preservation memory is combined
21 with a mirrored or backup copy of data as it existed at a selected point in time to roll the
22 data of the mass storage device back to a time prior to the creation of the mirrored or
23 backup copy. As described above in reference to Figure 4, data blocks stored in a
24 preservation memory prior to being overwritten in the mass storage device can be used to

1 incrementally roll back the data in the mass storage device. Thus, if the preservation
2 memory includes a sequence of data blocks that were overwritten in the mass storage
3 device prior to a subsequent mirrored or backup copy of the data of the mass storage
4 device, the sequence of data blocks in the preservation memory can be used to
5 incrementally advance the mirrored or backup copy back in time as desired.

6 With overwritten data blocks stored in the preservation memory, the invention
7 can be used to establish a virtual mass storage device ("virtual device") that permits data
8 that existed previously in the mass storage device to be accessed. The virtual device 50
9 of Figure 5 appears, from the standpoint of the user or an application (i.e., data access
10 program 60) that accesses the virtual device, to contain the data that existed at a previous
11 point in time. For example, the virtual device 50 can be accessed using an operating
12 system and an associated file system as if the virtual device were another hard drive
13 installed on computer 10 of Figure 1.

14 One example of the manner in which the virtual device 50 can be used to access
15 data as it existed at a previous point in time is illustrated in Figure 5, using the same set
16 of data that has been previously described above in reference to Figures 3 and 4. In the
17 example of Figure 5, it has been determined that a valid set of data can be obtained by
18 rolling the data blocks 20e stored in the mass storage device 12 to the state in which they
19 existed at $T_{1.1}$ using the data blocks 40e stored in preservation memory 14, as has been
20 described in reference to Figure 4.

21 When virtual device 50 is used to access data, data access program 60 issues read
22 requests 70 to virtual device 50 rather than addressing the requests specifically to mass
23 storage device 12 or preservation memory 14. In this example, the read requests are used
24 to access the most recent valid set of data that existed prior to the data corruption event

1 that introduced corrupted data blocks D_x and D_y to the set of data blocks in mass storage
2 device 12. As described previously, it can be determined that a set of valid data that
3 existed just prior to the introduction of corrupted data block D_x can be obtained. In other
4 words, the valid set of data at $T_{1,1}$, which is just prior to the data corruption event of $T_{1,2}$,
5 can be obtained.

6 Upon receiving a request for data, as it existed in the previous, non-corrupted
7 state, the virtual device determines whether the data satisfying the request is to be
8 obtained from mass storage device 12 or preservation memory 14. If the read request is
9 directed to a data block having a non-corrupted version that has been stored in
10 preservation memory 14 at or after the data corruption event (i.e., at or after $T_{1,2}$), the
11 oldest such data block is accessed to respond to the read request. For example, if data
12 access program 60 requests a data block at the "D" position that is associated with the
13 most recent set of valid data, the read request is processed by accessing the oldest non-
14 corrupted data block stored in the preservation memory at or after $T_{1,2}$. Thus, in response
15 to the request for the data block "D", the virtual device accesses data block D_1 , shown in
16 crosshatch in Figure 5. In a similar manner, a request for the "B" data block is processed
17 by accessing data block B stored in preservation memory 14, which is also shown in
18 crosshatch in Figure 5.

19 If, however, the read request is directed to a data block that does not have a non-
20 corrupted version stored in preservation memory 14 at or after the data corruption event
21 (i.e., at or after $T_{1,2}$), the corresponding data block is accessed from mass storage device
22 12 to respond to the read request. For instance, requests for data blocks at positions "A",
23 "C", and "E" are processed by accessing the corresponding data blocks A_1 , C, and E from
24 mass storage device 12, which are shown in crosshatch in Figure 5. In this way, data

1 access program 60 can request the full set of data blocks that existed at the previous, non-
2 corrupted state from virtual device 50 and receive, in response thereto, the set of data
3 blocks (A₁, B, C, D₁, and E). Virtual device 50 identifies which data blocks to obtain and
4 whether to obtain the data blocks from mass storage device 12 or preservation memory
5 14. Using the data blocks, data access program 60 does not need to know the details of
6 mass storage device 12, preservation memory 14, or the data blocks stored thereon, but
7 instead simply issues a request to virtual device 50 as if virtual device 50 contained the
8 prior set of data.

9 In one embodiment, the prior set of data can be reconstructed in the manner set
10 forth above in reference to Figure 5. Specifically, the prior set of non-corrupted data can
11 be obtained by establishing a virtual device and reading the prior set of non-corrupted
12 data as set forth above. Data access program 60 then makes the prior set of non-
13 corrupted data available as needed for substantially any use.

14 The embodiments of the present invention may comprise a special-purpose or
15 general-purpose computer that includes various components, as discussed in greater
16 detail below. Embodiments within the scope of the present invention may also include
17 computer-readable media for carrying or having computer-executable instructions or data
18 structures stored thereon. Such computer-readable media can be any available media that
19 can be accessed by a general-purpose or special-purpose computer. By way of example,
20 and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM,
21 CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage
22 devices, or any other medium which can be used to carry or store desired program code
23 means in the form of computer-executable instructions or data structures and which can
24 be accessed by a general-purpose or special-purpose computer.

1 When information is transferred or provided over a network or another
2 communications connection (either hardwired, wireless, or a combination of hardwired or
3 wireless) to a computer, the computer properly views the connection as a computer-
4 readable medium. Thus, any such connection is properly termed a computer-readable
5 medium. Combinations of the above should also be included within the scope of
6 computer-readable media. Computer-executable instructions comprise, for example,
7 instructions and data which may cause a general-purpose computer, special-purpose
8 computer, or special-purpose processing device to perform a certain function or group of
9 functions.

10 Figure 6 and the following discussion are intended to provide a brief, general
11 description of a suitable computing environment in which the invention may be
12 implemented. Although not required, the invention can be described in the general
13 context of computer-executable instructions, such as program modules, being executed
14 by computers in network environments. Generally, program modules include routines,
15 programs, objects, components, data structures, etc. that perform particular tasks or
16 implement particular abstract data types. Computer-executable instructions, associated
17 data structures, and program modules represent examples of the program code means for
18 executing steps of the methods disclosed herein. The particular sequence of such
19 executable instructions or associated data structures represents examples of
20 corresponding acts for implementing the functions described in such steps.

21 Those skilled in the art will appreciate that the invention may be practiced in
22 network computing environments with many types of computer system configurations,
23 including personal computers, hand-held devices, mobile telephones, personal digital
24 assistants ("PDAs"), multi-processor systems, microprocessor-based or programmable

1 consumer electronics, network PCs, minicomputers, mainframe computers, and the like,
2 one example having been presented in Figure 1. The invention may also be practiced in
3 distributed computing environments where local and remote processing devices are
4 linked (either by hardwired links, wireless links, or by a combination of hardwired or
5 wireless links) through a communications network and both the local and remote
6 processing devices perform tasks.

7 With reference to Figure 6, an example system for implementing the invention
8 includes a general-purpose computing device in the form of a computer 120, including a
9 processing unit 121, a system memory 122, and a system bus 123 that couples various
10 system components including system memory 122 to processing unit 121. Computer 120
11 and the associated component illustrated in Figure 6 represent a more detailed illustration
12 of computer 10 of Figure 1. System bus 123 may be any of several types of bus
13 structures including a memory bus or memory controller, a peripheral bus, and a local bus
14 using any of a variety of bus architectures. System memory may include read only
15 memory ("ROM") 124 and random access memory ("RAM") 125. A basic input/output
16 system ("BIOS") 126, containing the basic routines that help transfer information
17 between elements within the computer 120, such as during start-up, may be stored in
18 ROM 124.

19 Computer 120 may also include a magnetic hard disk drive 127 for reading from
20 and writing to a magnetic hard disk 139, a magnetic disk drive 128 for reading from or
21 writing to a removable magnetic disk 129, and an optical disk drive 130 for reading from
22 or writing to removable optical disk 131 such as a CD-ROM or other optical media.
23 Magnetic hard disk drive 127, magnetic disk drive 128, and optical disk drive 130 are
24 connected to system bus 123 by a hard disk drive interface 132, a magnetic disk drive-

1 interface 133, and an optical drive interface 134, respectively. The drives and their
2 associated computer-readable media provide nonvolatile storage of computer-executable
3 instructions, data structures, program modules and other data for computer 120.
4 Although the example environment described herein employs a magnetic hard disk 139, a
5 removable magnetic disk 129 and a removable optical disk 131, other types of computer
6 readable media for storing data can be used, including magnetic cassettes, flash memory
7 cards, digital versatile disks, Bernoulli cartridges, RAMs, ROMs, and the like.

8 Program code means comprising one or more program modules may be stored on
9 hard disk 139, magnetic disk 129, optical disk 131, ROM 124, or RAM 125, including an
10 operating system 135, one or more application programs 136, other program modules
11 137, and program data 138. A user may enter commands and information into computer
12 120 through keyboard 140, pointing device 142, or other input devices (not shown), such
13 as a microphone, joy stick, game pad, satellite dish, scanner, or the like. These and other
14 input devices are often connected to processing unit 121 through a serial port interface
15 146 coupled to system bus 123. Alternatively, the input devices may be connected by
16 other interfaces, such as a parallel port, a game port or a universal serial bus ("USB"). A
17 monitor 147 or another display device is also connected to system bus 123 via an
18 interface, such as video adapter 148. In addition to the monitor, personal computers
19 typically include other peripheral output devices (not shown), such as speakers and
20 printers.

21 Computer 120 may operate in a networked environment using logical connections
22 to one or more remote computers, such as remote computers 149a and 149b. Remote
23 computers 149a and 149b may each be another personal computer, a server, a router, a
24 network PC, a peer device or other common network node. Remote computers 149a and

1 149b may include many or all of the elements described above relative to the computer
2 120, although only memory storage devices 150a and 150b and their associated
3 application programs 136a and 136b are illustrated in Figure 6. The logical connections
4 depicted in Figure 6 include a local area network (LAN) 151 and a wide area network
5 (WAN) 152 that are presented here by way of example and not limitation. Such
6 networking environments are commonplace in office-wide or enterprise-wide computer
7 networks, intranets and the Internet.

8 When used in a LAN networking environment, computer 120 is connected to
9 local network 151 through a network interface adapter 153 or similar adapter. When
10 used in a WAN networking environment, computer 120 may include a modem 154, a
11 wireless link, or other means for establishing communications over the wide area network
12 152, such as the Internet. Modem 154, which may be internal or external, is connected to
13 system bus 123 via the serial port interface 146. In a networked environment, program
14 modules depicted relative to computer 120, or portions thereof, may be stored in the
15 remote memory storage device. It will be appreciated that the network connections
16 shown are exemplary and other means of establishing communications over wide area
17 network 152 may be used.

18 The present invention may also be embodied in other specific forms without
19 departing from its spirit or essential characteristics. The described embodiments are to be
20 considered in all respects only as illustrative and not restrictive. The scope of the
21 invention is, therefore, indicated by the appended claims rather than by the foregoing
22 description. All changes, which come within the meaning and range of equivalency of
23 the claims, are to be embraced within their scope.

24 What is claimed and desired to be secured by United States Letters Patent is: